Pipelined Iterative Solvers with Kernel Fusion for Graphics Processing Units

KARL RUPP, JOSEF WEINBUB, ANSGAR JÜNGEL, and TIBOR GRASSER,
Vienna University of Technology

arXiv:1410.4054v2 [cs.MS] 15 Dec 2014

We revisit the implementation of iterative solvers on discrete graphics processing units and demonstrate the benefit of implementations using extensive kernel fusion for pipelined formulations over conventional implementations of classical formulations. The proposed implementations with both CUDA and OpenCL are freely available in ViennaCL and are shown to be competitive with or even superior to other solver packages for graphics processing units. Highest performance gains are obtained for small to medium-sized systems, while our implementations are on par with vendor-tuned implementations for very large systems. Our results are especially beneficial for transient problems, where many small to medium-sized systems instead of a single big system need to be solved.

Categories and Subject Descriptors: G.1.3 [**Numerical Analysis**]: Numerical Linear Algebra

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Iterative Solvers, Conjugate Gradient Method, BiCGStab Method, GMRES Method, GPU, OpenCL, CUDA

ACM Reference Format:
Karl Rupp, Josef Weinbub, Ansgar Jüngel, and Tibor Grasser, 2015. Pipelined iterative solvers with kernel fusion for graphics processing units. *ACM Trans. Math. Softw.* X, Y, Article 00 (December 201X), 24 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

The need for the solution of a linear system of equations described by a sparse matrix A and a right hand side vector b is ubiquitous in computational science and engineering. Most prominently, discretizations of linear partial differential equations by means of the finite element or the finite volume method directly lead to such systems. Smaller-sized systems may be solved using sparse direct solvers, whereas iterative solvers are preferred or even necessary for large systems, eventually supplemented by preconditioning techniques of various degrees of sophistication.

The fine-grained parallelism of iterative solvers from the family of Krylov methods is particularly attractive for massively parallel hardware such as graphics processing units (GPUs), whereas much more effort is required to expose the parallelism in sparse direct solvers appropriately [Kim and Eijkhout 2013; Schenk et al. 2008].

This work is supported by the Austrian Science Fund (FWF) through the grants P23296 and P23598. The authors also thank AMD and NVIDIA for hardware donations and Joachim Schöberl for providing access to a system equipped with two NVIDIA K20m for benchmarking purposes.
Author's addresses: Karl Rupp, Josef Weinbub and Tibor Grasser, Institute for Microelectronics, Vienna University of Technology, Gußhausstraße 27-29/E360, A-1040 Wien, Austria; email: {rupp, weinbub, grasser}@iue.tuwien.ac.at; Ansgar Jüngel, Institute for Analysis and Scientific Computing, Vienna University of Technology, Wiedner Hauptstraße 8-10/E101, A-1040 Wien, Austria; email: juengel@asc.tuwien.ac.at
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 201X ACM 0098-3500/201X/12-ART00 $15.00
DOI: http://dx.doi.org/10.1145/0000000.0000000

Sparse matrix-vector products - essential parts of Krylov methods - have been studied in detail for GPUs [Ashari et al. 2014; Baskaran and Bordawekar 2008; Bell and Garland 2009; Greathouse and Daga 2014] and for INTEL's many-integrated core (MIC) architecture [Liu et al. 2013; Saule et al. 2014], based on which a unified format also well-suited for multi-core processors has been proposed recently [Kreutzer et al. 2014]. Similarly, vendor-tuned implementations of the vector operations required in addition to the sparse matrix-vector products for implementing sparse iterative solvers from the family of Krylov methods are available. A disadvantage of current accelerators is their connection to the host system via the PCI-Express bus, which is often a bottleneck both in terms of latency as well as bandwidth. This mandates a certain minimum system size to amortize the overhead of data transfer through the PCI-Express bus in order to obtain any performance gains over an execution on the host.

Two programming models are currently in widespread use for general purpose computations on GPUs: CUDA is a proprietary programming model for NVIDIA GPUs [Nickolls et al. 2008] providing its own compiler wrapper, whereas OpenCL is a royalty-free open standard maintained by the Khronos Group[1] and is typically provided as a shared library. Although OpenCL can also be used for NVIDIA GPUs, the richer CUDA toolchain has resulted in a higher share of research on general purpose computations on GPUs using CUDA. Also, slight performance differences of CUDA and OpenCL, caused by different degrees of compiler optimizations or differences in the implementation rather than through differences in the programming model, have been reported [Fang et al. 2011; Karimi et al. 2010]. Automated translators such as Swan [Harvey and De Fabritiis 2011] or CU2CL [Martinez et al. 2011] have been developed to reduce the maintenance effort of CUDA and OpenCL branches. However, only a subset of CUDA and OpenCL is supported by these translators, limiting their applicability particularly for highly optimized implementations. Consequently, portable software libraries targeting GPUs are currently driven into providing support for both CUDA and OpenCL, for example PARALUTION[2], VexCL[3], or ViennaCL[4].

A substantial amount of research has been conducted on various preconditioning techniques for iterative solvers on GPUs including algebraic multigrid [Bell et al. 2012; Gandham et al. 2014; Richter et al. 2014; Wagner et al. 2012], incomplete factorizations [Li and Saad 2013; Naumov 2012], or sparse approximate inverses [Dehnavi et al. 2013; Lukash et al. 2012; Sawyer et al. 2012]. Nevertheless, hardware-efficient and scalable black-box preconditioners for GPUs are not available, but instead the use of problem-specific information is required [Yokota et al. 2011]. Taking preconditioner setup costs into account, iterative solvers using simple diagonal preconditioners or no preconditioner at all are often observed to be competitive in terms of time-to-solution for small to mid-sized systems, where e.g. the asymptotic optimality of multigrid preconditioners is not yet dominant [Wagner et al. 2012]. Similarly, matrix-free methods cannot be used with complicated black-box preconditioners in general.

In this work we consider three popular iterative solvers: The conjugate gradient (CG) method for symmetric positive definite systems [Hestenes and Stiefel 1952], the stabilized bi-conjugate gradient (BiCGStab) method for non-symmetric positive definite systems [van der Vorst 1992], and the generalized minimum residual (GMRES) method for general systems [Saad and Schultz 1986]. In contrast to previous work with a focus on the optimization of sparse matrix-vector products [Ashari et al. 2014; Baskaran and Bordawekar 2008; Bell and Garland 2009; Greathouse and Daga 2014; Kreutzer et al.

[1]Khronos Group, OpenCL: http://www.khronos.org/opencl/
[2]PARALUTION library: http://www.paralution.com/
[3]VexCL library: https://github.com/ddemidov/vexcl/
[4]ViennaCL library: http://viennacl.sourceforge.net/

2014; Liu et al. 2013; Saule et al. 2014], we consider the optimization potential of the full solvers rather than restricting the optimization to a single kernel. After a careful evaluation of the limiting resources for different system sizes and different densities of nonzeros in the system matrix, pipelining and kernel fusion techniques are presented in Section 2 to resolve these bottlenecks to the extent possible. The key principle in pipelined techniques is to apply not only a single operation to a datum loaded from main memory, but to chain multiple operations together to reduce the overall number of loads and stores to global memory. Pipelining is typically achieved by fusing multiple compute kernels, but compute kernels may also be fused only to reduce the overall number of kernel launches, not exhibiting any pipelining effect. Pipelining and kernel fusion are then applied to the CG method, the BiCGStab method, and the GMRES method in Section 3, leading to more efficient solver implementations than those using a sequence of calls to the basic linear algebra subprograms (BLAS) in vendor-tuned libraries. Section 4 then compares the proposed solver implementations with existing solver implementations for GPUs available in the software libraries CUSP[5], MAGMA[6], and PARALUTION[7], demonstrating a substantial performance gain of pipelined implementations for small systems without sacrificing performance for large systems. Our benchmark results demonstrate the benefit of kernel fusion and pipelining techniques for GPUs from AMD and NVIDIA for the CG method, the BiCGStab method, and the GMRES method, and clearly outline that these techniques have not been applied extensively in the context of GPU computing before.

The obtained execution times are also compared with those obtained from CPU-based implementations in the PETSc[8] library to demonstrate that CPU-based implementations are superior for typical sparse systems below about $3\,000$ unknowns. Our results, similar to previous investigations [Lee et al. 2010], also falsify wide-spread misconceptions of extreme performance gains using GPUs. We show that performance gains of GPUs over power-equivalent dual-socket CPU machines are below an order of magnitude on average. This holds true also for large problem sizes and when initial data setup costs on GPUs are not taken into account. Finally, Section 5 discusses the implications of our findings to software design and the need for more tightly integrated future hardware generations.

2. IMPLEMENTATION TECHNIQUES FOR FAST ITERATIVE SOLVERS

The purpose of this section is to identify the general bottlenecks of the typical building blocks of iterative solvers and to present techniques for mitigating their detrimental effects on performance. A schematic view of a machine (*host*) equipped with a discrete GPU connected via PCI-Express is given in Fig. 1, where the following key features are schematically depicted using a terminology similar to OpenCL:

— Threads are collected in *workgroups*, where each workgroup provides dedicated memory shared across threads in the workgroup. Thread synchronizations within a workgroup are possible inside a compute kernel, but a global synchronization of all workgroups is typically only possible by launching a new kernel. Although global synchronization primitives and spin locks through atomic operations are used occasionally, these techniques are not sufficiently portable across different hardware and thus not further considered.

[5]CUSP library: http://cusplibrary.github.io/
[6]MAGMA library: http://icl.cs.utk.edu/magma/
[7]PARALUTION library: http://www.paralution.com/
[8]PETSc library: http://www.mcs.anl.gov/petsc/

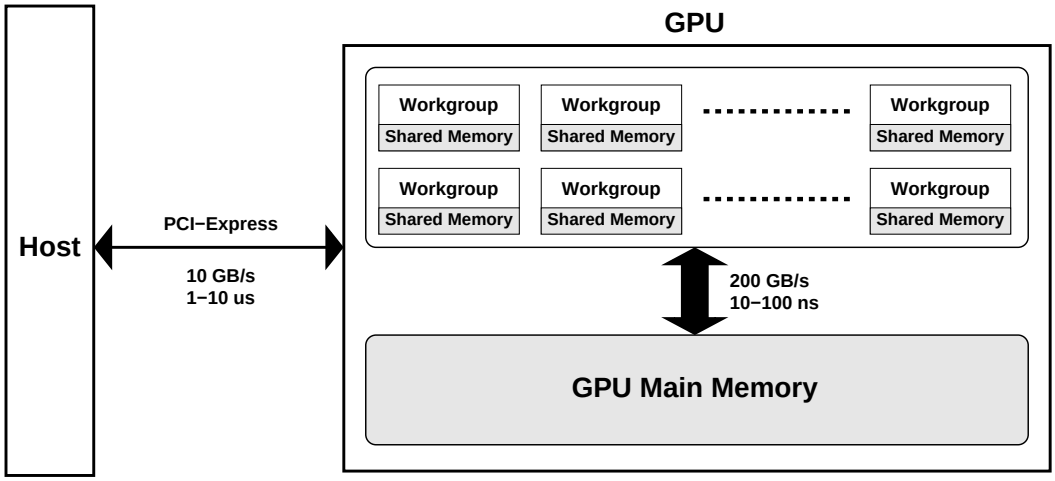

Fig. 1. Schematic view of a GPU board connected to the host via PCI-Express at a bandwidth of about 10 GB/sec and a latency on the order of 10 microseconds. Each workgroup of threads can be synchronized through shared memory, but global synchronization is available only through separate kernel launches.

— If a kernel launch is initiated on the host, it takes at least a few microseconds until the kernel will launch on a GPU. This is because a kernel launch on the GPU requires a message from the host to trigger the execution, entailing high latency for communication across PCI-Express. This latency of kernel launches can be hidden if another kernel is currently active on the GPU, in which case the PCI-Express message for launching the new kernel is received asynchronously.
— Memory access latency of GPU main memory is around three orders of magnitude smaller than the latency of messages across the PCI-Express bus.
— The memory bandwidth between GPU main memory and the GPU compute units can be more than ten times higher than the bandwidth of the PCI-Express bus connecting host and GPU. Current high-end GPUs offer over 200 GB/sec memory bandwidth, whereas the current PCI-Express 3.0 offers up to 15.75 GB/sec for a 16-lane slot.

The remainder of this section quantifies the overhead of PCI-Express latency and presents techniques for reducing the number of kernel launches to reduce the detrimental latency effect.

2.1. PCI-Express Latency

At the very least, iterative solvers executed on the GPU need to communicate information about the current residual norm to the host. In the typical case of a communication of the residual norm in each iteration for convergence checks, the time required for a data transfer from the device to the host represents a lower bound for the time required for an iterative solver iteration. An OpenCL benchmark for PCI-Express data communication shown for an NVIDIA Tesla C2050 in Fig. 2 exhibits a latency-dominated regime for message sizes below ten kilobytes, where the transfer time is around eight microseconds. Latency-dominated data transfer from the device to the host takes almost twice as long, because a transfer initiation from the host is required first. Similar timings and bandwidths are obtained on other GPUs both with PCI-Express 2.0 and 3.0. Our overall observation in Section 4 is that NVIDIA GPUs show slightly lower latency than AMD GPUs on the Linux-based machines used for the comparison.

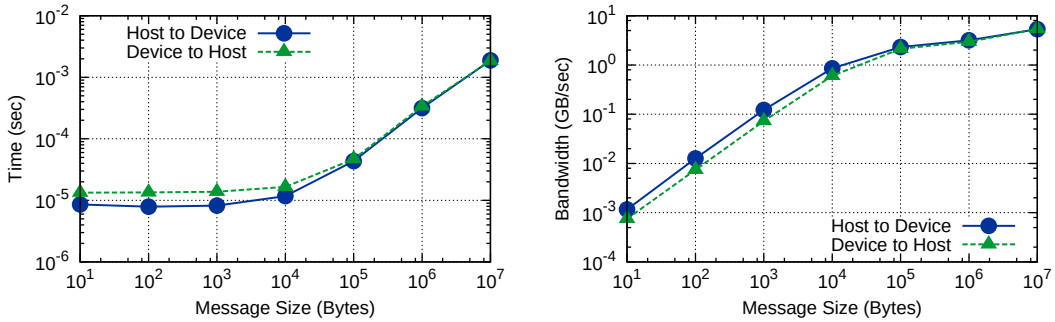

Fig. 2. Plot of median values for execution time (left) and obtained bandwidth (right) from 100 host-device data transfers over PCI-Express 2.0 using an NVIDIA Tesla C2050. The benchmark uses the OpenCL functions clEnqueueWriteBuffer() and clEnqueueReadBuffer() in a blocking manner so that the respective function only returns after the data is sent or received. Message sizes below 10^4 bytes are limited by latency, not PCI-Express bandwidth.

To better understand the latency induced by PCI-Express transfer, consider a high-end GPU with 200 GB/s memory bandwidth. Within the PCI-Express latency of 8 microseconds, the GPU can load or store 1.6 megabytes of data from main memory assuming full saturation of the memory channel, which amounts to 200 000 values in double precision and which we will refer to as *latency barrier*. Consequently, GPUs suffer from inherent performance constraints for any kernel limited by memory bandwidth whenever the total amount of data loaded or stored is significantly below the latency barrier. On the other hand, many practical applications induce systems with storage requirements for the unknowns close to or even below the latency barrier. In such case, iterative solver implementations for GPUs need to keep the latency-induced overhead as small as possible by packing multiple operations into each kernel.

2.2. Kernel Fusion

As a prototypical example for many iterative solvers, consider the sequence of operations

$$q = Ap \tag{1}$$
$$\alpha = \langle p, q \rangle \tag{2}$$

for a scalar value α, vectors p and q, and a sparse square matrix A. Conventional implementations based on BLAS routines involve the following steps:

(1) Call the sparse matrix-vector product kernel for computing (1). For a standard compressed sparse row (CSR) representation of the sparse matrix, a typical OpenCL kernel body is as follows (cf. [Baskaran and Bordawekar 2008; Bell and Garland 2009]):

```
for (uint i = get_global_id(0); i < size; i += get_global_size(0))
{
  double q_at_i = 0;
  for (uint j = A_row[i]; j < A_row[i+1]; ++j)
    q_at_i += A_values[j] * p[A_col[j]];
  q[i] = q_at_i;
}
```

where A_row and A_col are the arrays holding the row and column indices in the CSR storage format, and A_values holds the nonzero entries. High-performance implementations of sparse matrix-vector products for the CSR format are more involved than the simple example shown here [Ashari et al. 2014; Greathouse and Daga 2014].

(2) Compute the partial results of $\langle p, q \rangle$ for the subvectors assigned to each of the thread workgroups.
(3) If α is stored on the GPU, launch another kernel using a single thread workgroup to sum the partial results. If α is stored on the host, transfer the partial results to the host and perform the summation there.

Although this conventional implementation can reuse vendor-tuned routines, the multiple kernel launches are detrimental to performance for data sizes below the PCI-Express latency barrier.

On closer inspection, the operations (1) and (2) can be computed more efficiently by fusing compute kernels: Since the respective values in q and p are already available in the GPU processing elements when computing the matrix-vector product, they can be reused to compute the partial results for each thread workgroup of the inner product. The fused kernel body for the CSR format is as follows:

```
//
// Part 1: Matrix-vector product
//
double p_in_q = 0;
for (uint i = get_global_id(0); i < size; i += get_global_size(0))
{
  double q_at_i = 0;
  for (uint j = A_row[i]; j < A_row[i+1]; ++j)
    q_at_i += A_values[j] * p[A_col[j]];
  q[i]    = q_at_i;
  p_in_q += q_at_i * p[i];      // extra operation for <p, q>
}

//
// Part 2: Reduction to obtain contribution from thread workgroups:
//
__local double shared_buf[BUFFER_SIZE];
shared_buffer[get_local_id(0)] = p_in_q;
for (uint stride=get_local_size(0)/2; stride > 0; stride /= 2)
{
  barrier(CLK_LOCAL_MEM_FENCE);
  if (get_local_id(0) < stride)
    shared_buf[get_local_id(0)] += shared_buf[get_local_id(0) + stride];
}

if (get_local_id(0) == 0)
  partial_result[get_group_id(0)] = shared_buf[0];
```

First, the matrix-vector kernel from the previous snippet is only slightly augmented to accumulate the partial results for each thread in `p_in_q`. Extra logic could be employed to explicitly avoid reading `p[i]` from global memory if the respective diagonal entry of A is nonzero, but `p[i]` may still be available in cache anyway. Then, a reduction using a shared memory buffer `shared_buf` of appropriate size `BUFFER_SIZE` is applied to obtain the sum over all threads within a thread workgroup. Finally, the first thread in each thread workgroup writes the partial result of the workgroup to a temporary buffer `partial_result`. The summation of the values in `partial_result` is carried out on the host as outlined in the third step above.

A comparison of execution times of the conventional implementation with the implementation using the fused kernel is given in Fig. 3. In both cases the final reduction step for the partial results from 128 thread workgroups has been computed on the host and is included in the timings. Two types of matrices have been compared: The first family of matrices with four randomly (with a uniform distribution over all column indices) distributed nonzeros per row is limited by latency for systems with up to 10^4 unknowns. A performance gain of about 20 percent is obtained from the use of a fused kernel, which reduces the number of kernels required from two to just one. At sys-

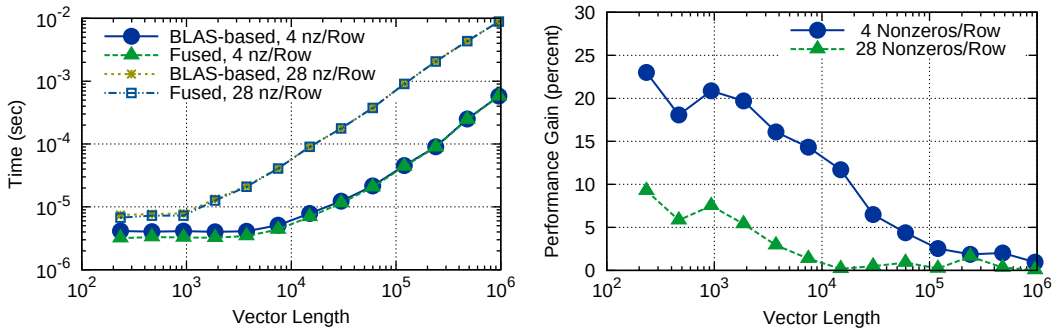

Fig. 3. Total time required to run the operations (1) and (2) for different matrix and vector sizes on an NVIDIA Tesla C2050. If the vector size is below $10\,000$ entries, the total time is dominated by the latency for enqueuing the kernel, not the kernel execution time.

tem sizes above 10^5 unknowns, a performance gain of a few percent is still obtained because the vector q does not need to be reloaded from memory when computing the inner product (2). The second matrix type with 28 randomly distributed nonzeros per row is limited by the kernel execution time for system sizes below 10^3 unknowns. This is because each thread needs to process 28 nonzeros per row in A, which results in a larger execution time than the pure PCI-Express latency. Nevertheless, a performance gain of up to ten percent is obtained for smaller systems, yet there is no notable performance gain or loss at larger system sizes due to diminishing savings from reusing values from q for computing the inner product.

It is not only possible to fuse the first stage in the inner product computation with the matrix-vector product, but one can also fuse the second stage (summation of partial results) with subsequent operations. Since the summation result is usually needed in each thread workgroup, the final summation has to be computed in each thread workgroup in such case. These redundant computations are usually well below the PCI-Express latency barrier and thus faster than the use of a separate host-device transfer or a dedicated summation kernel. While kernel fusion can in principle be applied to an arbitrary number of vector updates, the global synchronization points induced by matrix-vector products as well as inner products are natural boundaries for fusing compute kernels. However, not every inner product induces a separate synchronization point: The partial summation stage of several inner products may also be computed within the same kernel, which is then followed by a second kernel computing the final results of the inner products and possibly other vector operations.

3. PIPELINED ITERATIVE METHODS FOR GRAPHICS PROCESSING UNITS

The implementation of the CG method, the BiCGStab method, and the GMRES method are investigated in depth in the following. Each of these solvers is analyzed for the number of kernel launches to evaluate latency. The kernel fusion techniques outlined in Section 2 are applied to reduce the number of kernel launches whenever appropriate. We restrict our investigations to the execution on a single GPU, as this is the most frequent use case. Nevertheless, the optimizations applied in this section can also be transferred to a multi-GPU setting, where additional data exchange between GPUs via the PCI-Express bus entails similar cost. This allows, for example, to pack multiple partial results from inner products into a single memory buffer transfer to minimize latency.

3.1. Conjugate Gradient Method

Several variations of the classical CG method [Hestenes and Stiefel 1952] have been proposed in the past, cf. [Barkai et al. 1985; Chronopoulos and Gear 1989; Ghysels and Vanroose 2014]. Also, techniques for merging multiple solver iterations have been proposed, but they do not find broad acceptance in practice because of numerical instabilities [Saad 1985]. In the following, the classical CG method and a pipelined version are compared, where the latter has already been developed for vector machines [Chronopoulos and Gear 1989], revisited for extreme-scale scalability [Ghysels and Vanroose 2014], and implemented in field-programmable gate arrays [Strzodka and Göddeke 2006]:

	Algorithm 1: Classical CG		**Algorithm 2:** Pipelined CG
1	Choose x_0	1	Choose x_0
2	$p_0 = r_0 = b - Ax_0$	2	$p_0 = r_0 = b - Ax_0$
3		3	Compute and store Ap_0
4		4	$\alpha_0 = \langle r_0, r_0\rangle / \langle p_0, Ap_0\rangle$
5		5	$\beta_0 = \alpha_0^2 \langle Ap_0, Ap_0\rangle / \langle r_0, r_0\rangle - 1$
6	**for** $i = 0$ *to convergence* **do**	6	**for** $i = 1$ *to convergence* **do**
7	Compute and store Ap_i	7	
8	Compute $\langle p_i, Ap_i\rangle$	8	
9	$\alpha_i = \langle r_i, r_i\rangle / \langle p_i, Ap_i\rangle$	9	
10	$x_{i+1} = x_i + \alpha_i p_i$	10	$x_i = x_{i-1} + \alpha_{i-1} p_{i-1}$
11	$r_{i+1} = r_i - \alpha_i Ap_i$	11	$r_i = r_{i-1} - \alpha_{i-1} Ap_{i-1}$
12		12	$p_i = r_i + \beta_{i-1} p_{i-1}$
13		13	Compute and store Ap_i
14		14	Compute $\langle Ap_i, Ap_i\rangle$, $\langle p_i, Ap_i\rangle$
15	Compute $\langle r_{i+1}, r_{i+1}\rangle$	15	Compute $\langle r_i, r_i\rangle$
16		16	$\alpha_i = (r_i, r_i\rangle / (p_i, Ap_i\rangle$
17	$\beta_i = \langle r_{i+1}, r_{i+1}\rangle / \langle r_i, r_i\rangle$	17	$\beta_i = \alpha_i^2 \langle Ap_i, Ap_i\rangle / \langle r_i, r_i\rangle - 1$
18	$p_{i+1} = r_{i+1} + \beta_i p_i$	18	
19	**end**	19	**end**

A direct implementation of Algorithm 1 using one call to a matrix-vector product routine and five calls to BLAS routines per solver iteration is straightforward. Optimizations of the matrix-vector products on lines 2 and 7 in Algorithm 1 and lines 2, 3 and 13 in Algorithm 2 have been investigated in detail for different matrix formats on GPUs in the past [Baskaran and Bordawekar 2008; Bell and Garland 2009]. The inner products in lines 8 and 15 in the classical CG formulation impose synchronization by either splitting the operation into two kernels or by requiring a host-device transfer. In particular, the residual norm computed in line 15 is typically required on the host for convergence checks. The vectors r and p are loaded in lines 10 and 11, but have to be reloaded for the search vector update operation in line 18.

The pipelined version in Algorithm 2 is based on the relation

$$\langle r_{i+1}, r_{i+1}\rangle = \frac{\alpha_i^2 \langle Ap_i, Ap_i\rangle - \langle r_i, r_i\rangle}{\langle r_i, r_i\rangle} = \alpha_i^2 \langle Ap_i, Ap_i\rangle / \langle r_i, r_i\rangle - 1 \tag{3}$$

to compute p_i in line 18 of Algorithm 1 without having computed $\langle r_{i+1}, r_{i+1}\rangle$ yet. We note that it has been stated in the literature that precomputing inner products involving the vectors p_i and r_i by using recursion formulas based only on inner products of p_j and r_j with $j < i$ may lead to unstable algorithms [Chronopoulos and Gear 1989; Saad 1985]. However, the computation of β_i involves Ap_i, resulting in a stable algorithm based on experiences from multiple groups in different application contexts [Barkai et al. 1985; Chronopoulos and Gear 1989; Strzodka and Göddeke 2006].

Relation (3) allows for a rearrangement of Algorithm 1 such that all vector updates can be computed right after each other (lines 10, 11, and 12 in Algorithm 2). An application of kernel fusion not only allows for computing all three vector updates within a single kernel instead of three, but also for avoiding a reload of p_{i-1} and r_i (line 12) when using registers for intermediate values. Furthermore, all three inner products in Algorithm 2 can be computed simultaneously, allowing all intermediate results to be communicated to the host with a single data transfer. More precisely, the first reduction stage for the inner products $\langle Ap_i, Ap_i\rangle$, $\langle p_i, Ap_i\rangle$, and $\langle r_i, r_i\rangle$ can be computed within the same kernel sharing the same buffer for intermediate results. Then, the second reduction stage for obtaining the final results is either computed by only a single additional kernel launch, or by communicating all partial results with a single transfer to the host, which performs the summation. The data size of a single partial result is about one kilobyte, hence the data transfer time remains in the latency-dominated regime even if three of them are packed together, cf. Fig. 2.

To further enhance data reuse, we fuse the matrix-vector product in line 13 with the inner products in line 14, so that the result values of the matrix vector product can be processed right before they are written to GPU RAM. Thus, the computation of $\langle Ap_i, Ap_i\rangle$ and $\langle p_i, Ap_i\rangle$ comes at reduced data transfer cost, because the j-entry of Ap_i has just been computed, and the j-th entry of p_i may still be available in cache. Similarly, the inner product $\langle r_i, r_i\rangle$ in line 15 is fused with the vector update kernel for lines 10, 11, and 12.

In summary, we propose the following implementation of Algorithm 2:

— Compute lines 10, 11, and 12 in one kernel and store the reduction results of each workgroup for the computation of $\langle r_i, r_i\rangle$ in line 15 in a temporary buffer.
— Compute lines 13 and 14 in one kernel and append the reduction results of each workgroup for the computation of the inner products in line 14 to the same temporary buffer.
— Communicate the temporary buffer to the host, where the final reduction is computed to obtain $\langle r_i, r_i\rangle$, α_i, and β_i from lines 15, 16 and 17.

Since $\langle r_i, r_i\rangle$ is available for monitoring the residual norm on the host, a convergence check can be applied in each iteration with no extra effort. The proposed implementation requires only two kernel launches per iteration and one host-device data transfer. In contrast, a direct translation of the classical CG algorithm into BLAS-routines requires at least six kernel launches (lines 7, 8, 10, 11, 15, and 18) and may involve a second host-device data transfer for $\langle p_i, Ap_i\rangle$. Consequently, we expect an up to threefold performance gain for small systems in the latency-dominated regime. Because p_i and r_i do not need to be loaded from memory twice per iteration, a performance gain of a few percent may also be obtained for large systems with very few nonzeros per row.

3.2. BiCGStab

BiCGStab is an attractive iterative solver for systems described by nonsymmetric matrices, because the transposed operator A^{T} is not required. Based on the initial derivation [van der Vorst 1992], a pipelined method with only two global synchronizations has been proposed [Jacques et al. 1999]. Later, a variant with only a single global synchronization has been proposed at the cost of an application of the transposed operator in the setup stage [Yang and Brent 2002]. Also, a preconditioned BiCGStab method overlapping global communication with the application of the preconditioner has been developed [Krasnopolsky 2010]. A preliminary optimization study of the classical BiCGStab for GPUs is also available [Anzt et al. 2014], for which we postpone a comparison to Section 4.

In analogy to the classical BiCGStab algorithm, the pipelined BiCGStab implementation proposed in this work does not require the transposed operator A^{T} to be available and is similar to the one proposed with two global synchronizations [Jacques et al. 1999]. A comparison with the classical BiCGStab algorithm [Saad 2003] is as follows:

	Algorithm 3: Classical BiCGStab		**Algorithm 4:** Pipelined BiCGStab
1	Choose x_0	1	Choose x_0
2	$p_0 = r_0 = b - Ax_0$	2	$p_0 = r_0 = b - Ax_0$
3	Choose r_0^* arbitrary	3	Choose r_0^* arbitrary
4	Compute $\langle r_0, r_0^* \rangle$	4	Compute $\langle r_0, r_0^* \rangle$
5	**for** $i = 0$ *to convergence* **do**	5	**for** $i = 0$ *to convergence* **do**
6	Compute and store Ap_i	6	Compute and store Ap_i
7	Compute $\langle Ap_i, r_0^* \rangle$	7	Compute $\langle Ap_i, r_0^* \rangle$
8	$\alpha_i = \langle r_i, r_0^* \rangle / \langle Ap_i, r_0^* \rangle$	8	$\alpha_i = \langle r_i, r_0^* \rangle / \langle Ap_i, r_0^* \rangle$
9	$s_i = r_i - \alpha_i Ap_i$	9	$s_i = r_i - \alpha_i Ap_i$
10	Compute and store As_i	10	Compute and store As_i
11	Compute $\langle As_i, s_i \rangle$, $\langle As_i, As_i \rangle$	11	Compute $\langle As_i, s_i \rangle$, $\langle As_i, As_i \rangle$
12		12	Compute $\langle As_i, r_0^* \rangle$
13		13	$\beta_i = -\frac{\langle As_i, r_0^* \rangle}{\langle Ap_i, r_0^* \rangle}$
14	$\omega_i = \langle As_i, s_i \rangle / \langle As_i, As_i \rangle$	14	$\omega_i = \langle As_i, s_i \rangle / \langle As_i, As_i \rangle$
15	$x_{i+1} = x_i + \alpha_i p_i + \omega_i s_i$	15	$x_{i+1} = x_i + \alpha_i p_i + \omega_i s_i$
16	$r_{i+1} = s_i - \omega_i As_i$	16	$r_{i+1} = s_i - \omega_i As_i$
17	Compute $\langle r_{i+1}, r_0^* \rangle$	17	
18	$\beta_i = \frac{\langle r_{i+1}, r_0^* \rangle}{\langle r_i, r_0^* \rangle} \times \frac{\alpha_i}{\omega_i}$	18	
19	$p_{i+1} = r_{i+1} + \beta_i (p_i - \omega_i A_i)$	19	$p_{i+1} = r_{i+1} + \beta_i (p_i - \omega_i A_i)$
20		20	Compute $\langle r_{i+1}, r_0^* \rangle$
21	**end**	21	**end**

The classical BiCGStab method in Algorithm 3 requires a global synchronization after line 7 to compute α_i for use in line 8. Similarly, synchronizations are also required after line 11 to compute ω_i for use in line 14 and after line 17 to compute β for use in line 18. In analogy to the classical CG method, the search direction vector p_{i+1} (line 19) cannot be updated together with the approximated solution x_{i+1} (line 15) and the residual vector r_{i+1} (line 16). Consequently, additional loads from GPU main memory are required. Overall, two calls to routines for sparse matrix-vector products and at least eight calls to BLAS level 1 routines are needed in a conventional implementation of the classical BiCGStab method. Four host-device data transfers are required if each inner product induces a data transfer between host and device. An additional call to a BLAS level 1 routine and a host-device transfer are necessary if the residual norm is recomputed explicitly in each iteration.

The pipelined BiCGStab version in Algorithm 4 allows for improved data reuse by shifting the calculation of β_i to line 13 through

$$\beta_i = \frac{\langle r_{i+1}, r_0^* \rangle}{\langle r_i, r_0^* \rangle} \times \frac{\alpha_i}{\omega_i} = \frac{\langle s_i - \omega_i As_i, r_0^* \rangle}{\langle r_i, r_0^* \rangle} \times \frac{\langle r_i, r_0^* \rangle}{\omega_i \langle Ap_i, r_0^* \rangle} = \frac{\langle s_i, r_0^* \rangle}{\omega_i \langle Ap_i, r_0^* \rangle} - \frac{\langle As_i, r_0^* \rangle}{\langle Ap_i, r_0^* \rangle} .$$

Using the orthogonality $\langle s_i, r_0^* \rangle = \langle r_i - \alpha_i Ap_i, r_0^* \rangle = 0$ one arrives at

$$\beta_i = -\frac{\langle As_i, r_0^* \rangle}{\langle Ap_i, r_0^* \rangle} ,$$

which we found to be numerically stable based on our experiments. This derivation of a pipelined BiCGStab version is similar to the modification of the classical CG method in Algorithm 1 to obtain the pipelined Algorithm 2. The minor price to pay for this rearrangement is the calculation of $\langle As_i, r_0^* \rangle$ in line 12.

The next step is to apply kernel fusion extensively to the pipelined BiCGStab version in Algorithm 4. The calculation of $\langle As_i, r_0^* \rangle$ can be fused with the sparse matrix product in line 10 together with the calculation of $\langle As_i, s_i \rangle$ and $\langle As_i, As_i \rangle$ in line 11. Similarly, lines 6 and 7 are fused to a single kernel computing a matrix-vector product and the first reduction stage of the inner product. The vector update in line 9 is fused with the second reduction stages for the inner products needed to compute α_i in line 8. Since the residual norm is obtained via

$$\langle r_{i+1}, r_{i+1} \rangle = \langle s_i, s_i \rangle - 2\omega_i \langle s_i, As_i \rangle + \omega_i^2 \langle As_i, As_i \rangle$$

for which $\langle s_i, As_i \rangle$ and $\langle As_i, As_i \rangle$ are computed in line 11 and needed for the calculation of ω_i in line 14, we augment the update kernel for the computation of s_i in line 9 with the first reduction stage for $\langle s_i, s_i \rangle$. The partial results are transferred to the host together with the partial results for all other inner products after line 12, where β_i and ω_i are computed. Finally, the vector updates in lines 15, 16 and 19 as well as the first reduction stage for the inner product in line 20 are fused into another kernel.

Overall, the proposed pipelined BiCGStab implementation of Algorithm 4 consists of four kernel launches and one host-device transfer of the partial results from the four inner products $\langle As_i, s_i \rangle$, $\langle As_i, As_i \rangle$, $\langle As_i, r_0^* \rangle$, and $\langle s_i, s_i \rangle$:

— Compute the matrix-vector product in line 6 and the partial results for the two inner products required for α_i in line 8.
— Compute s_i in line 9 by redundantly computing α_i in each thread workgroup from the partial results of the inner products $\langle r_i, r_0^* \rangle$ and $\langle Ap_i, r_0^* \rangle$.
— Compute and store As_i (line 10) and the partial results for the inner products in lines 11 and 12.
— Communicate all partial results for the inner products to the host, sum them there and perform a convergence check.
— Compute the vector updates in lines 15, 16, and 19 as well as the partial results for the inner product in line 20.

In comparison, the BiCGStab implementation proposed in [Anzt et al. 2014] requires five kernel launches and three reductions, while a BLAS-based implementation of the classical method requires at least eight kernel launches and four additional kernel launches or host-device transfers for the second reduction stage in the computation of the inner products. Therefore, a moderate improvement over the pipelined implementation in [Anzt et al. 2014] and a two- to three-fold performance gain over purely BLAS-based implementations in the latency-dominated regime are expected, assuming that kernel launches and host-device transfers entail comparable latency.

3.3. GMRES

In contrast to the CG and BiCGStab methods, the GMRES method requires to store the full Krylov basis rather than only the current search direction vector, leading to an increase in the number of operations with each iteration [Saad and Schultz 1986]. To limit the computational expense, the GMRES method is typically restarted after m iterations, which is denoted by GMRES(m). Typical values for m are in the range of 20 to 50. Smaller values tend to slow down the overall convergence, whereas higher values increase the computational cost and may lead to more time spent in the orthogonalization rather than the matrix-vector product, making GMRES less attractive when compared to other methods.

Three methods for the computation of an orthonormal Krylov basis from a set of linearly independent vectors $\{v_k\}_{k=1}^m$ are common [Saad 2003], further algorithms employed for the orthogonalization in a multi-GPU setting with significantly different constraints in terms of communication can be found in [Yamazaki et al. 2014]:

— *Classical Gram-Schmidt*: The k-th vector of the basis is obtained as

$$w_k \leftarrow \sum_{i=1}^{k-1} \langle v_i, v_k \rangle v_i \ , v_k \leftarrow v_k - w_k$$

followed by a normalization of v_k. The inner products $\langle v_i, v_k \rangle$ are independent and can be computed in parallel.

— *Modified Gram-Schmidt*: An accumulation of round-off errors in the basis vectors v_k may lead to a loss of orthogonality as the basis is augmented. Higher numerical robustness than for the classical Gram-Schmidt method has been observed for

$$v_k \leftarrow v_k - \langle v_i, v_k \rangle v_i$$

for i from 1 to $k-1$ rather than forming a single update vector w_k. The disadvantage of the modified Gram-Schmidt method is the reduced parallelism: Instead of computing all inner products $\langle v_i, v_k \rangle$ concurrently, only one inner product can be computed at a time, followed by a vector update.

— *Householder reflections*: The Krylov basis may also be obtained through Householder reflections $P_k = (I - \beta_k u_k u_k^{\mathrm{T}})$ with identity matrix I, suitably chosen scalars β_k, and Householder vectors u_k. Similar to the modified Gram-Schmidt method, the Householder reflections have to be applied sequentially to obtain the Krylov basis. Although the method allows for the computation of an orthonormal basis up to machine precision, the method is less regularly used for implementations of the GMRES method due to its sequential nature.

In the following we consider the simpler GMRES method [Walker and Zhou 1994], which allows for a simpler solution of the minimization problem than the original formulation, but is otherwise comparable in terms of computational expense. A comparison of the restarted form and a pipelined formulation, both using the classical Gram-Schmidt method for higher efficiency on parallel architectures, is as follows:

Algorithm 5: Classical GMRES(m)

```
Choose x_0
r_0 = b − Ax_0
ρ_0 = ‖r_0‖_2
v_0 = r_0 = r_0/ρ_0
R_{i,j} = 0 for i, j ∈ {1, ..., m}
for i = 1 to m do
    v_i = A v_{i−1}
    for j = 1 to i − 1 do
        R_{j,i} = ⟨v_j, v_i⟩
    end
    for j = 1 to i − 1 do
        v_i = v_i − R_{j,i} v_j
    end
    v_i = v_i / ‖v_i‖
    ξ_i = ⟨r, v_i⟩
    r = r − ξ_i v_i
end
Solve Rη = (ξ_1, ..., ξ_m)
Update x_m = η_1 r + Σ_{i=2}^{m} η̃_i v_{i−1}
```

Algorithm 6: Pipelined GMRES(m)

```
Choose x_0
r_0 = b − Ax_0
ρ_0 = ‖r_0‖_2
v_0 = r_0 = r_0/ρ_0
R_{i,j} = 0 for i, j ∈ {1, ..., m}
for i = 1 to m do
    v_i = A v_{i−1}
    for j = 1 to i − 1 do
        R_{j,i} = ⟨v_j, v_i⟩
    end
    for j = 1 to i − 1 do
        v_i = v_i − R_{j,i} v_j
    end
    v_i = v_i / ‖v_i‖
    ξ_i = ⟨r, v_i⟩ (first stage)

end
for i = 1 to m do
    ξ_i = ⟨r, v_i⟩ (second stage)
end
Solve Rη = (ξ_1, ..., ξ_m)
Update x_m = η_1 r + Σ_{i=2}^{m} η_i v_{i−1}
```

with $\tilde{\eta_i} = \eta_i + \eta_1 \xi_{i-1}$ to account for the updates of the residual r.

The main difference between the classical formulation in Algorithm 5 and the pipelined formulation in Algorithm 6 involves the update of the residual vector in line 16 of Algorithm 5. Because of the orthonormality of $\{v_k\}_{k=1}^m$, the inner product in line 15 remains unchanged when using exact arithmetic. Similarly, since the values ξ_i do not enter the Gram-Schmidt process, the values in the matrix R remain unchanged so that round-off errors only affect the right hand side vector in line 21. Our numerical experiments indicate that round-off errors in ξ_i are dominated by round-off errors in the classical Gram-Schmidt process and therefore do not affect the overall numerical stability of the solver. Also, the convergence monitors proposed in [Walker and Zhou 1994] do not require updates of the residual and are based on the values ξ_i only. Therefore, the full convergence history is still accessible before solving the minimization problem in line 21, allowing for a correct handling of early convergence. Nevertheless, $m-1$ unnecessary steps of the Gram-Schmidt process will be carried out if convergence is obtained right at the first iteration, but this is rarely encountered for unpreconditioned solvers in practice.

The benefit of removing the residual update from the Gram-Schmidt orthogonalization is that extensive kernel fusion can be applied to obtain an implementation of Algorithm 6 with almost no host-device communication. To begin, the reduction stage of the inner products in line 9 can be computed in two ways: The first option is a specialized matrix-vector routine for tall matrices if all Krylov vectors are stored as either the rows or the columns of a matrix. The second option is to fuse multiple inner products into the same kernel if all Krylov vectors reside in distinct buffers [Rupp et al. 2013]. With both options the second reduction stage for computing $R_{j,i}$ in line 9 is fused with the vector updates in line 12 and also with the first reduction stage for computing $\|v_i\|$ needed in line 14. The normalization of v_i in line 14 is carried out by a kernel first computing the second reduction stage for $\|v_i\|$, then scaling v_i and directly computing the first reduction stage for the obtaining ξ_i in line 15. Consequently, no data transfer between host and device is required during the Gram-Schmidt orthogonalization. An asynchronous transfer of the intermediate values for ξ_i can be inserted at the end of each orthogonalization step for a better monitoring of the convergence process.

After the Gram-Schmidt process, the intermediate results for computing ξ_i are transferred to the host if not already transferred asynchronously, where the final values ξ_i are computed. Similarly, the triangular matrix R is transferred to the host. After the triangular system R is inverted, the result vector containing the values η_i is transferred to the device and the update of the result vector x_m is computed in line 22 using in a single kernel similar to the vector update in line 12.

Overall, the proposed implementation of the pipelined GMRES(m) method in Algorithm 6 requires two kernel launches in the first iteration and four kernel launches in subsequent iterations:

— Compute the matrix-vector product in line 7 and the first reduction stage for $\langle v_{i-1}, v_i\rangle$.
— Compute the first reduction stage for the inner products $\langle v_j, v_i\rangle$ in line 9 with j ranging from 1 to $i-2$.
— Compute the second reduction stage for the inner products $\langle v_j, v_i\rangle$ in line 9 for j from 1 to $i-1$, use the results directly for computing the vector update in 12, and compute the first reduction stage for $\|v_i\|$.
— Compute the second reduction stage for $\|v_i\|$, normalize v_i, and compute ξ_i.

A conventional implementation of the classical GMRES(m) method in Algorithm 5 requires at least seven kernel launches and may involve several host-device data exchanges per iteration. Thus, an up to two-fold performance gain in the latency-dominated regime is expected.

4. BENCHMARK RESULTS

The implementations proposed in this work are implemented in the 1.6.2 release of the free open-source linear algebra library ViennaCL[9] and are compared in the following with the implementations in the free open-source libraries CUSP[10] 0.4.0, MAGMA[11] 1.5.0 (linked with INTEL MKL 11.0), and PARALUTION[12] 0.8.0. Since CUSP and MAGMA are based on CUDA, benchmark data for AMD GPUs could only be obtained with ViennaCL and PARALUTION. All four libraries are used in an out-of-the-box manner without additional target-specific tuning in order to reflect typical use cases.

In addition to classical implementations of the CG, BiCGStab and GMRES methods, MAGMA also provides pipelined implementations of the CG method (four kernels with custom sparse matrix-vector product) and the BiCGStab method (nine kernels using vendor-tuned sparse matrix-vector products) [Anzt et al. 2014]. CUSP and PARALUTION implement classical formulations of all three iterative solvers in the comparison. Our numerical experiments showed that the implementations of the classical CG and BiCGStab methods in MAGMA show similar performance to PARALUTION. Therefore, we do not include timings for the classical implementations in MAGMA in our benchmark result plots, but instead only report execution times for the pipelined variants.

All tests were carried out on Linux-based machines running the CUDA 6.0 SDK on NVIDIA GPUs with GPU driver version 331.20 and the AMD APP SDK 2.9 with GPU driver version 13.352.1014 on AMD GPUs. An NVIDIA Tesla C2050, an NVIDIA Tesla K20m, an AMD FirePro W9000, and an AMD FirePro W9100 were used for a comparison, representing the latest two generations of high-end workstation models from each vendor. Error-correcting code memory was disabled on all four GPUs for better comparison. Since all operations are limited by the available memory bandwidth, the obtained results are also representative for a broader range of mid- to high-end consumer GPUs with comparable memory bandwidth.

In addition to GPU-benchmarks, we also compare with the execution times obtained with the CPU-based PETSc[13] library on a dual-socket system equipped with INTEL Xeon E5-2620 CPUs, where parallel execution is based on the Message Passing Interface (MPI)[14] using MPICH[15] 3.1. The fastest execution time from runs with 1, 2, 4, and 8 MPI ranks for each system size are taken for comparison. However, it should be noted that a comparison with a CPU-based library needs to be interpreted with care, because our benchmarks only compare the time taken per solver iteration, not the time required for copying the data to the GPU or for obtaining the result vector.

Execution times per iterative solver iteration are computed from the median value of ten solver runs with a fixed number of 30 iterations for each solver. In our experiments we have not observed any significant differences in the number of solver iterations required for convergence of the classical implementation and the pipelined implementation, hence the execution time per solver iteration is a suitable metric for comparison.

[9] ViennaCL library: http://viennacl.sourceforge.net/
[10] CUSP library: http://cusplibrary.github.io/
[11] MAGMA library: http://icl.cs.utk.edu/magma/
[12] PARALUTION library: http://www.paralution.com/
[13] PETSc library: http://www.mcs.anl.gov/petsc/
[14] Message Passing Interface Forum: http://www.mpi-forum.org/
[15] MPICH library: http://www.mpich.org/

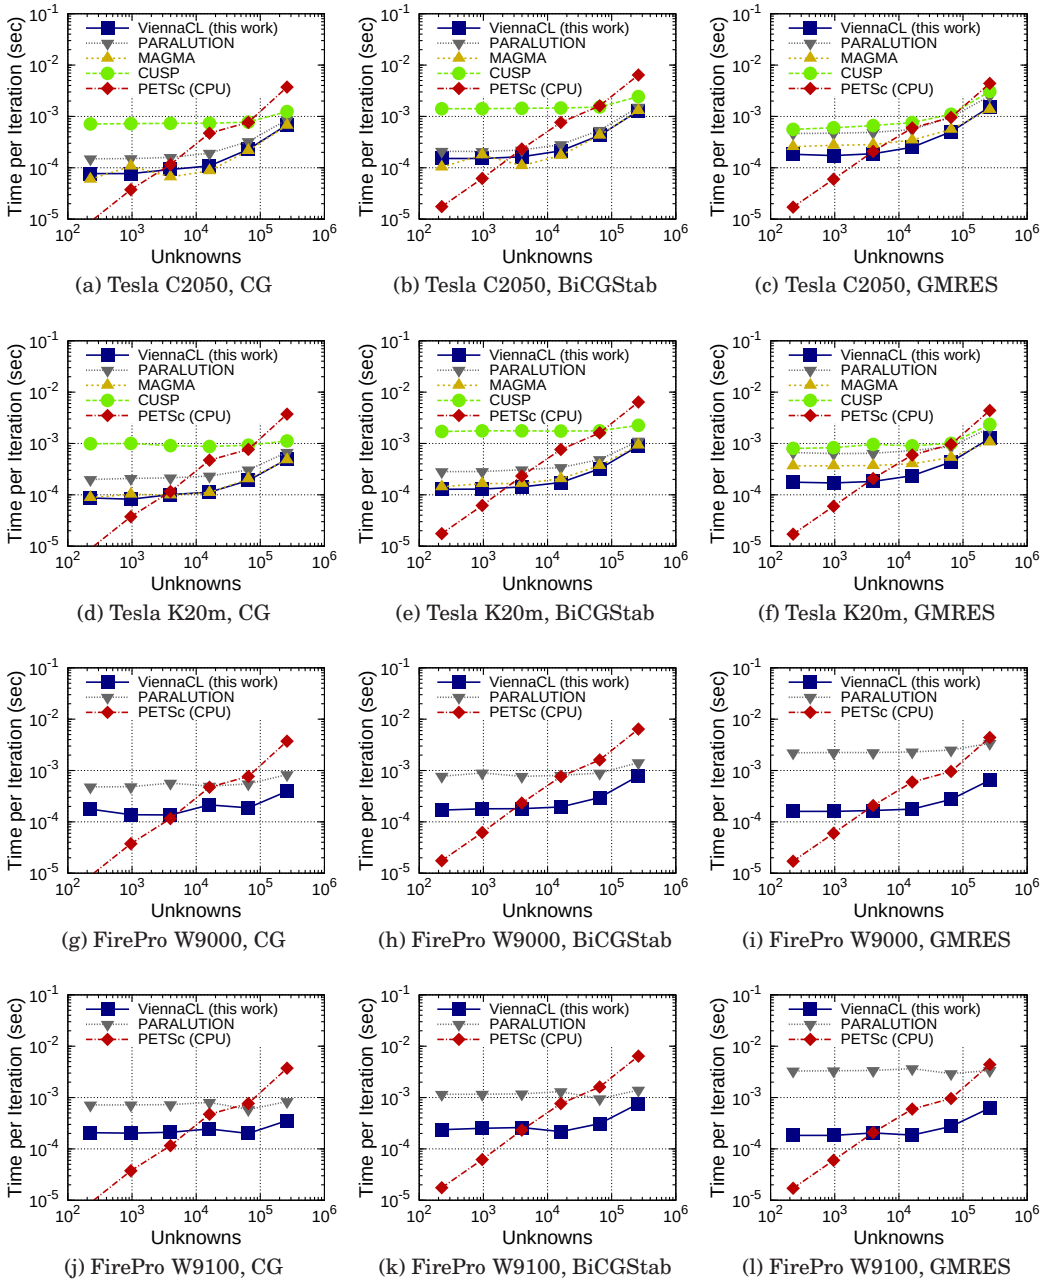

(a) Tesla C2050, CG (b) Tesla C2050, BiCGStab (c) Tesla C2050, GMRES

(d) Tesla K20m, CG (e) Tesla K20m, BiCGStab (f) Tesla K20m, GMRES

(g) FirePro W9000, CG (h) FirePro W9000, BiCGStab (i) FirePro W9000, GMRES

(j) FirePro W9100, CG (k) FirePro W9100, BiCGStab (l) FirePro W9100, GMRES

Fig. 4. Comparison of the solver time required per iteration for solving the Poisson equation using finite elements on triangular grids in two spatial dimensions. The proposed pipelined implementations in ViennaCL as well as the pipelined implementations in MAGMA outperform classical implementations in CUSP and PARALUTION for system sizes below 10^5 thanks to a smaller number of kernel launches and better data reuse.

4.1. Linear Finite Elements for the Poisson Equation in 2D

We consider the execution time obtained with linear finite elements applied to the solution of the Poisson equation on the unit rectangle on a hierarchy of uniformly refined unstructured triangular meshes as a first benchmark. The resulting systems consist of 225, 961, $3\,969$, $16\,129$, $65\,025$ and $261\,121$ equations, respectively, and cover a broad range of typical system sizes solved on a single workstation. Results for CG, BiCGStab and GMRES using the ELLPACK sparse matrix format (cf. [Bell and Garland 2009] for a description) are given in Fig. 4 for the four GPUs considered in our comparison. Similar results are obtained for other matrix formats, because the execution times in this setting are primarily dominated by latency effects. The case of large system matrices, where kernel execution times are dominant, is considered in Section 4.3.

Results for the NVIDIA GPUs show that the pipelined methods for the CG and the BiCGStab methods in ViennaCL and MAGMA show the same performance for small systems. The small differences can be explained by the use of asynchronous (and hence non-deterministic) convergence checks in MAGMA, whereas ViennaCL uses synchronous checks. Although MAGMA's BiCGStab implementation uses five kernels rather than the proposed implementation with four kernels, no significant difference is visible in Fig. 4b and Fig. 4e. A comparison with the classical implementations in PARALUTION shows a roughly two-fold performance gain of pipelined implementations for small systems. The proposed pipelined GMRES implementation on NVIDIA GPUs is by about a factor of two faster in the latency-dominated regime than the implementation in MAGMA and about a factor of three faster than the implementations in CUSP and PARALUTION.

On AMD GPUs the differences between ViennaCL and PARALUTION are more pronounced, because PARALUTION cannot take advantage from some optimizations in CUBLAS for NVIDIA GPUs. Conversely, these results suggests that the CUDA runtime for NVIDIA GPUs is able to hide the overhead of kernel launches more efficiently. A three-fold difference in execution times is obtained for the CG method, which reflects the different number of kernel launches, namely two for the pipelined implementation and six for a conventional implementation. A four-fold difference in execution times is obtained for the BiCGStab method, again reflecting the reduction in the number of kernel launches and reduced host-device communication in the proposed pipelined implementation. The difference for GMRES is approximately ten-fold, because the Gram-Schmidt orthogonalization in PARALUTION calls one kernel per dot-product during the orthogonalization procedure.

Execution times for each solver iteration at system sizes below 10^4 are practically constant for both NVIDIA and AMD GPUs. Because this constant is about a factor of two larger for AMD GPUs and because the AMD GPUs in this comparison offer higher memory bandwidth, essentially constant execution times are obtained for systems with up to 10^5 unknowns for AMD GPUs. Only at system sizes above 10^5 unknowns, PCI-Express communication becomes negligible compared to kernel execution times, hence the performance of all libraries becomes similar and varies only mildly.

When comparing the execution times of GPU-based solvers with the execution times obtained with the CPU-based PETSc implementations, it is observed that the proposed pipelined implementations on GPUs are faster if systems carry more than about $3\,000$ unknowns on average. Depending on the underlying hardware and solver, up to $100\,000$ unknowns are needed with the conventional implementations in PARALUTION or CUSP to outperform the CPU-based implementations in PETSc. If initial data setup is taken into account, these cross-over points are shifted to even larger values, highlighting the importance of pipelining to increase the range of system sizes where GPU acceleration may pay off.

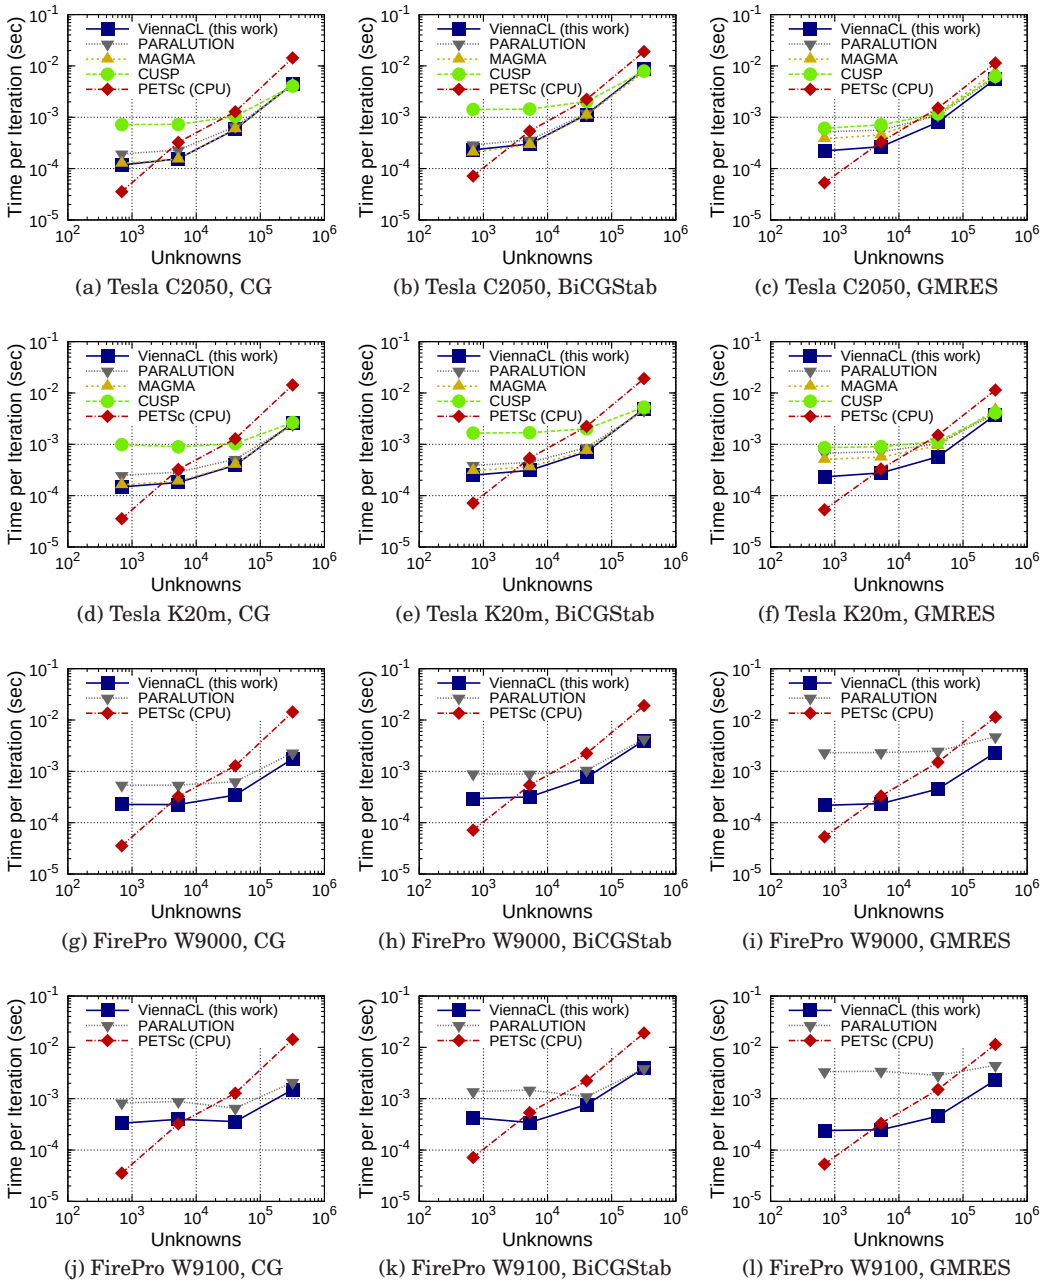

(a) Tesla C2050, CG (b) Tesla C2050, BiCGStab (c) Tesla C2050, GMRES

(d) Tesla K20m, CG (e) Tesla K20m, BiCGStab (f) Tesla K20m, GMRES

(g) FirePro W9000, CG (h) FirePro W9000, BiCGStab (i) FirePro W9000, GMRES

(j) FirePro W9100, CG (k) FirePro W9100, BiCGStab (l) FirePro W9100, GMRES

Fig. 5. Comparison of the solver time required per iteration for solving the linear elasticity model using finite elements in three spatial dimensions. The proposed pipelined implementations in ViennaCL as well as the pipelined implementations in MAGMA outperform other libraries for system sizes below 10^5 thanks to a smaller number of kernel launches and better data reuse.

Table I. Summary of symmetric and non-symmetric matrices taken from the Florida Sparse Matrix Collection for comparison. These matrices represent the set of real-valued, floating point square matrices used in earlier contributions on optimizing sparse matrix-vector products [Bell and Garland 2009; Kreutzer et al. 2014].

Name	Rows	Nonzeros	Nonzeros/Row	Symmetric
pdb1HYS	36 417	4 344 765	119.31	yes
cant	62 451	4 007 383	64.17	yes
consph	83 334	6 010 480	72.13	yes
shipsec1	140 874	7 813 404	55.46	yes
pwtk	217 918	11 643 424	53.39	yes
rma10	46 835	2 374 001	50.69	no
cop20k_A	121 192	2 624 331	21.65	no
scircuit	170 998	958 936	5.61	no
mac_econ_fwd500	206 500	1 273 389	6.17	no
RM07R	381 689	37 464 962	98.16	no
Hamrle3	1 447 360	5 514 242	3.81	no
kkt_power	2 063 494	13 612 663	7.08	no

4.2. Linear Finite Elements for Linear Elasticity in 3D

The second benchmark compares the execution time obtained with linear finite elements for numerical solutions of the linear elasticity model in three spatial dimensions. A hierarchy of uniformly refined tetrahedral meshes of the unit cube was used, resulting in system sizes of 693, $5\,265$, $40\,725$, and $319\,725$, respectively. Compared to the first benchmark, the average number of unknowns per row increases from about 7 to 60 for the largest system, resulting in a higher share of the execution time being spent on sparse matrix-vector products.

The results on NVIDIA GPUs in Fig. 5 show a similar trend as the results in Fig. 4: For small matrix sizes, the pipelined implementations of the CG and the BiCGStab methods in ViennaCL and MAGMA show similar performance. A two-fold performance gain over PARALUTION is obtained for the smallest system when using the CG method, which quickly diminish at larger system sizes due to more time spent on sparse matrix-vector products. While CUSP is about five times slower in the latency-limited regime for BiCGStab, the implementation in PARALUTION is less than a factor of two slower, suggesting that the CUDA-runtime is able to hide kernel launch latencies as well as host-device communication fairly well. Similar to the previous benchmark, the performance gain of the proposed pipelined implementation of GMRES is two-fold over MAGMA and three-fold over CUSP and PARALUTION.

Performance differences between ViennaCL and PARALUTION on the AMD GPUs are about two- to three-fold for the CG and BiCGStab method. For GMRES, a ten-fold performance advantage of the proposed pipelined implementation in the latency-dominated regime is obtained.

Although the system matrix carries more nonzeros than in the first benchmark, about 2×10^4 unknowns on NVIDIA GPUs and 10^5 unknowns on AMD GPUs are required such that kernel execution times hide performance penalties due to PCI-Express communication.

4.3. Florida Sparse Matrix Collection

The performance of the proposed pipelined implementations is compared in the following for matrices from the Florida Sparse Matrix Collection[16] used for the evaluation of sparse matrix-vector products in the past [Bell and Garland 2009; Kreutzer et al. 2014]. While the focus in the previous section was on demonstrating the benefit of the proposed implementations for small to medium-sized systems, the purpose of this

[16] http://www.cise.ufl.edu/research/sparse/matrices/

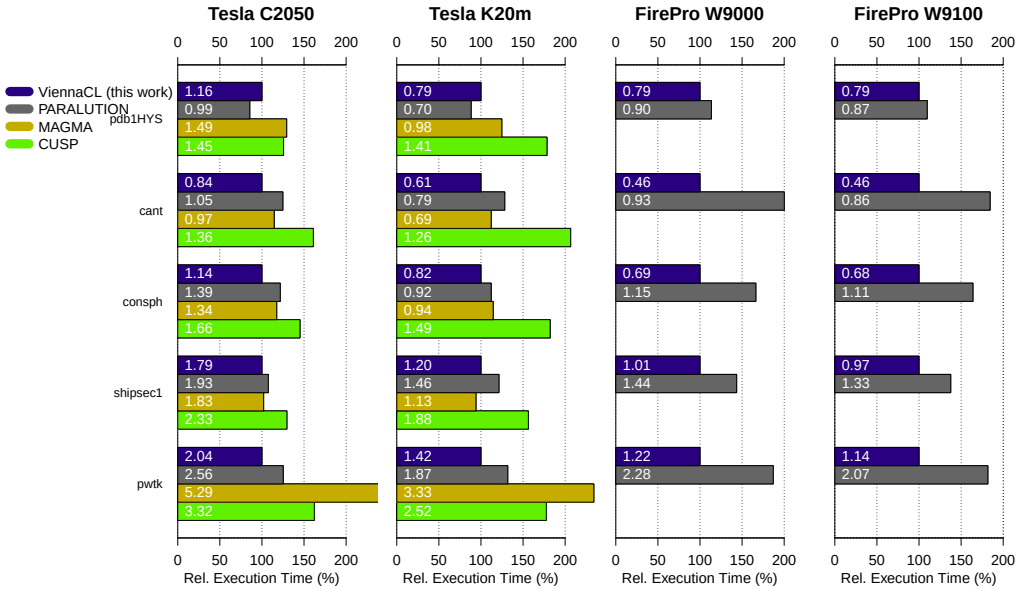

Fig. 6. Comparison of execution times per CG solver iteration for different systems from the Florida Sparse Matrix Collection relative to the proposed pipelined implementations. Absolute execution times in milliseconds are given inside each bar. ViennaCL implements the pipelined methods proposed in this work, MAGMA uses a similar pipelined implementation without using vendor-tuned kernels.

section is to show that the proposed implementations are also competitive for large systems. Thus, the implementations in PARALUTION and MAGMA (for BiCGStab and GMRES) are a-priori expected to provide the best performance, since they use the vendor-tuned sparse matrix-vector product kernels from NVIDIA's CUSPARSE library. In contrast, our implementations in ViennaCL rely on fused kernels, while CUSP implements the classical methods using its own set of sparse matrix-vector product kernels [Bell and Garland 2009].

Since OpenCL does not support complex arithmetic natively, we restrict our benchmark to real-valued matrices listed in Tab. I. The symmetric, positive definite matrices are used for benchmarking the implementations of the CG method, while the non-symmetric matrices are used for benchmarking the implementations of the BiCGStab and GMRES methods. All sparse matrix formats available in the respective library are compared using implementations in CUDA and, if available, OpenCL. The fastest combination is then taken for the comparison, since such a procedure resembles the typical user who picks the fastest sparse matrix format and the programming model with best performance for a particular application.

The benchmark results for the CG method in Fig. 6 show that the proposed solver implementation provides the best overall performance on all four devices. Although the pipelined CG method implemented in MAGMA is similar to the one proposed here, the performance difference reflects the importance of providing fast fused kernels. Similarly, the difference is particularly pronounced on AMD GPUs, where the performance of our proposed implementation is up to twice as high as the performance of PARALUTION, which needs to rely on its own kernels rather than using vendor-tuned implementations. A comparison of absolute execution times also shows that the AMD GPUs provide a better overall performance due to their higher memory bandwidth.

The comparison of execution times for the BiCGStab method in Fig. 7 shows similar performance of ViennaCL, PARALUTION, and MAGMA for NVIDIA GPUs on aver-

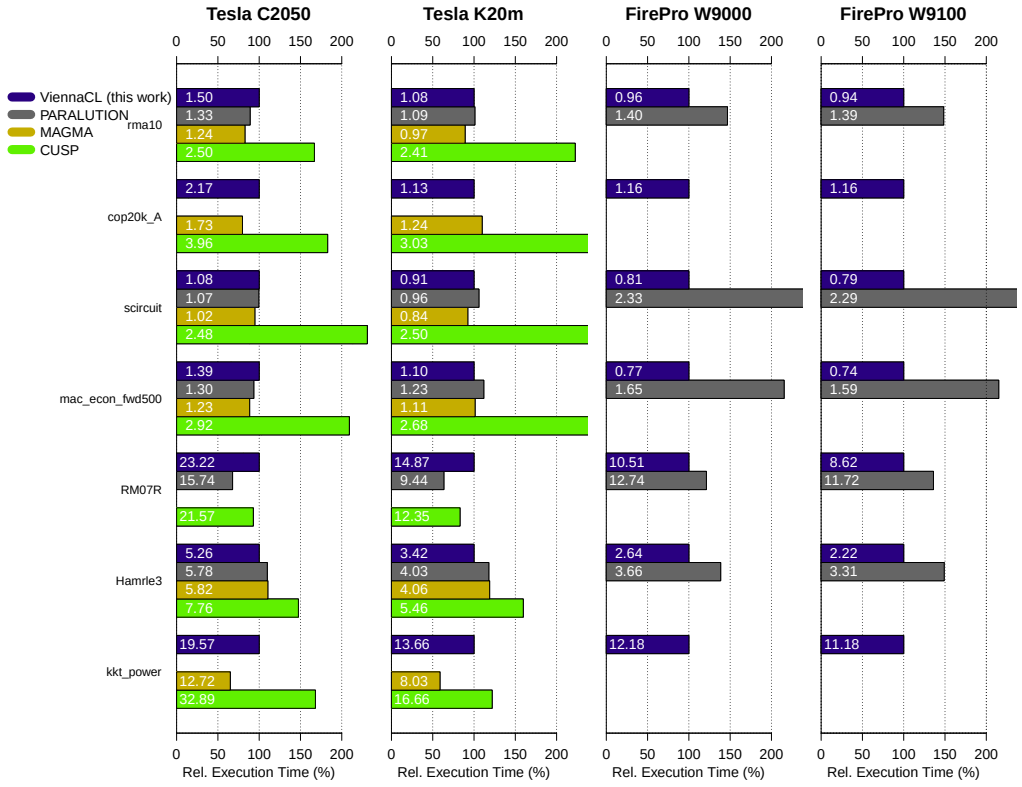

Fig. 7. Comparison of execution times per BiCGStab solver iteration for different systems from the Florida Sparse Matrix Collection relative to the proposed pipelined implementations. Absolute execution times in milliseconds are given inside each bar. The `cop20k_A` and the `kkt_power` matrices could not be tested with PARALUTION due to segmentation faults. The `RM07R` matrix could not be run with MAGMA since it did not pass a check for positive definiteness.

age: Depending on the device and the matrix considered, either of the three is the best choice. Since the proposed implementations do not contain any device-specific or matrix-specific optimizations, further tuning may provide further performance gains, whereas the implementations in PARALUTION and MAGMA using vendor-tuned kernels do not have this option. The custom sparse matrix-vector product kernels in CUSP result in about 50 percent higher execution times on average. On AMD GPUs, the performance gain over PARALUTION is about 60 percent on average. Similar to the results of the benchmark of the CG method, slightly higher overall performance can be obtained on AMD GPUs because of their higher memory bandwidth.

The benchmark results obtained for the GMRES method are depicted in Fig. 8 and show the same trend as the results obtained when comparing the implementations of the BiCGStab method. Depending on the device and the matrix considered, either of ViennaCL, PARALUTION, and MAGMA may be the best choice. In particular, no performance penalty from using pipelined implementations for large systems can be observed.

Finally, execution times for the proposed implementations of the three iterative solvers using CUDA and OpenCL are compared in Fig. 9. In all cases, the obtained execution times of CUDA and OpenCL are within a few percent, which is a negligible difference in practice.

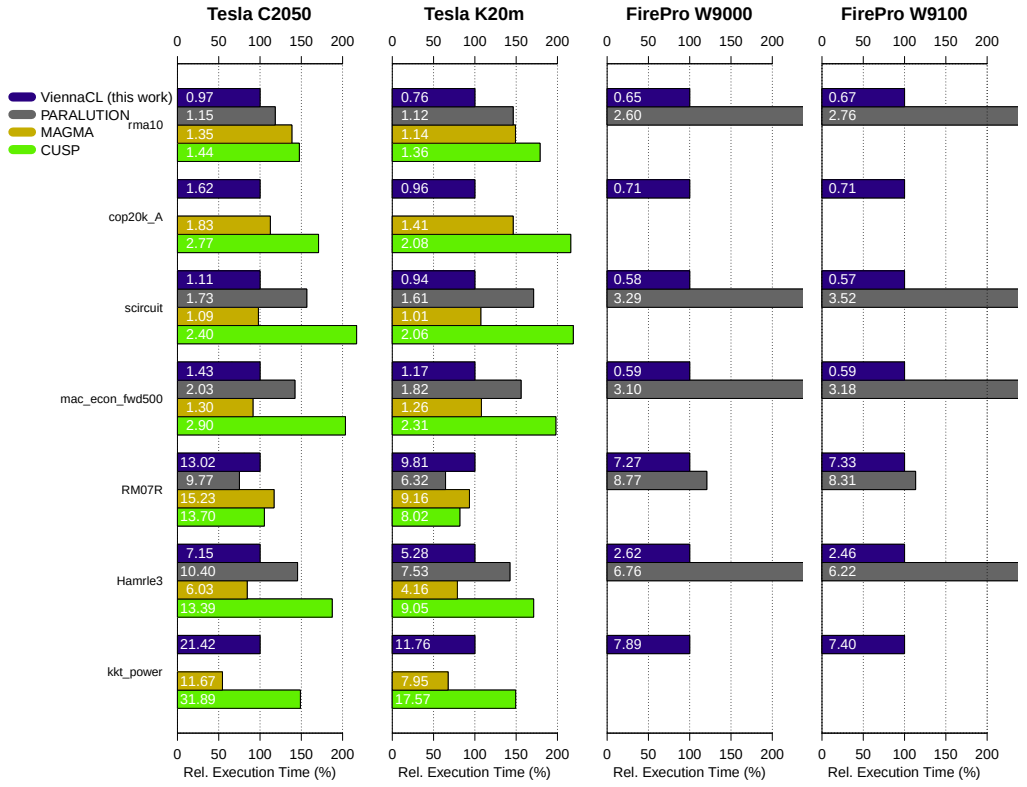

Fig. 8. Comparison of execution times per GMRES solver iteration for different systems from the Florida Sparse Matrix Collection relative to the proposed pipelined implementations. Absolute execution times in milliseconds are given inside each bar. The `cop20k_A` and the `kkt_power` matrices could not be tested with PARALUTION due to segmentation faults.

Overall, the benchmark results confirm that pipelined methods are not only favorable for smaller systems, where latency effects are significant, but also competitive for large systems. This is beneficial for code maintenance, as only a single implementation needs to be maintained. Furthermore, our results also suggest that a single implementation in OpenCL is sufficient, as the performance differences to CUDA are negligible. This, however, faces practical limitations, as user codes may be written only in CUDA and thus incompatible with OpenCL.

5. CONCLUSION

The proposed pipelined implementations of the CG, BiCGStab and GMRES methods address the latency-induced performance penalties of GPU-accelerated implementations for sparse systems with less than about 10^5 unknowns. Our comparison with other solver packages shows significant performance gains over conventional implementations for practically relevant problem sizes between 10^4 and 10^5 unknowns. A comparison for larger systems shows that the proposed implementations using fused kernels provide a performance competitive with implementations built on top of vendor-tuned kernels. As a consequence, our results suggest that future efforts on the optimization of compute kernels should not be restricted to standard BLAS or BLAS-like kernels, but additional performance can be obtained if also optimized implementations for fused kernels are provided. For example, not only the sparse matrix-vector

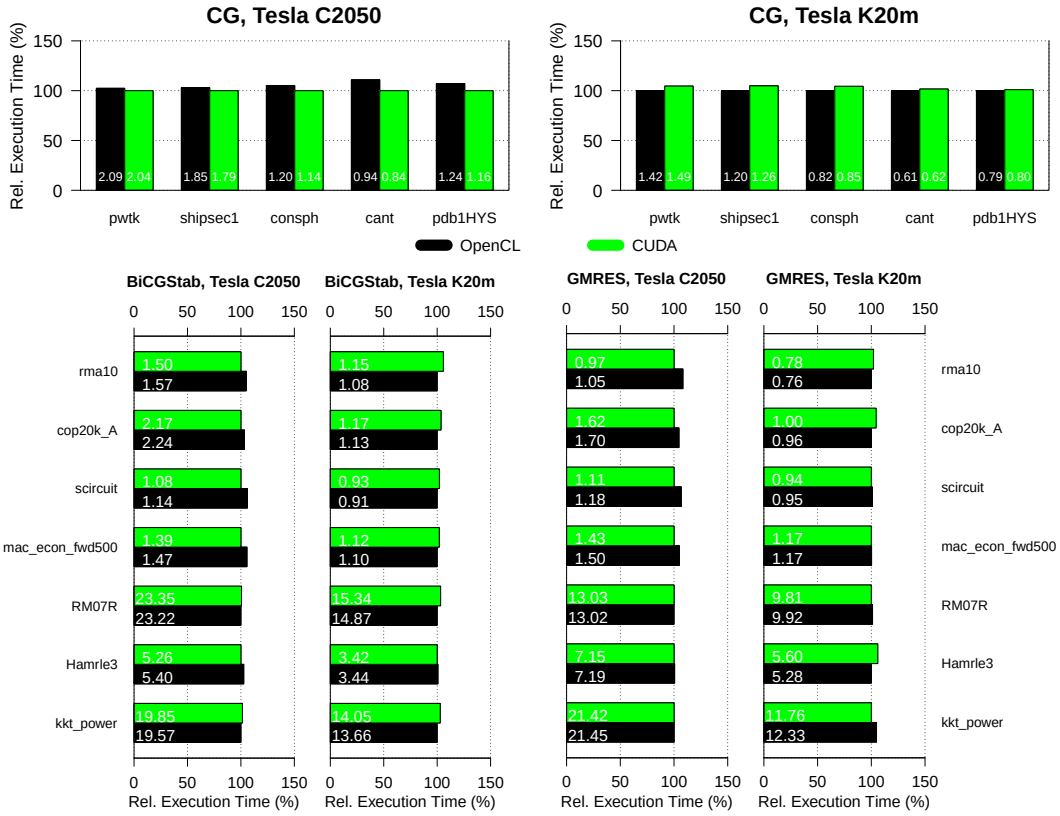

Fig. 9. Comparison of execution times obtained with CUDA and OpenCL for the CG method (top), the BiCGStab method (left), and the GMRES method (right). Relative execution times are with respect to the faster framework. Absolute execution times in milliseconds are given inside each bar. Overall, the performance differences of CUDA and OpenCL are negligible in practice, even though OpenCL shows slightly better performance overall.

product kernel, but also a kernel computing the sparse matrix-vector product plus the first reduction stage of inner products involving the result vector may offer superior performance for iterative solvers from the family of Krylov methods.

While an extensive use of pipelining and kernel fusion addresses latency issues and limited memory bandwidth, it also brings new challenges for the design of scientific software. To leverage the full potential of modern hardware, it is no longer sufficient to only use a fairly small set of vendor-tuned BLAS-kernels, but instead provide modular building blocks for minimizing communication of data.

Future GPUs as well as CPUs will see gains in memory bandwidth, but the latency induced by the PCI-Express bus will not change substantially. Therefore, the minimum system size required to get any performance gains on GPUs over CPUs will continue to grow. As a consequence, the replacement of the PCI-Express bus with a interconnect technology of lower latency is essential for making accelerators more attractive. Integrations of GPU units on the CPU die are one possible path to achieve lower latency. However, no benefit over a well-optimized, purely CPU-based implementations can be expected for the memory-bandwidth limited operations in iterative solvers, if both the accelerator and the CPU core share the same memory link.

The techniques applied in this work can also be extended to preconditioned iterative solvers. Not only can the application of the preconditioner be possibly fused with vector updates, but also the setup stage can benefit from fusing as many operations as possible into the same kernel. A rigorous application of these techniques to preconditioners is left for future work.

REFERENCES

H. Anzt, W. Sawyer, S. Tomov, P. Luszczek, I. Yamazaki, and J. Dongarra. 2014. Optimizing Krylov Subspace Solvers on Graphics Processing Units. In *IEEE Intl. Conf. Par. Dist. Sys. Workshops*. 941–949. DOI:http://dx.doi.org/10.1109/IPDPSW.2014.107

A. Ashari, N. Sedaghati, J. Eisenlohr, S. Parthasarathy, and P. Sadayappan. 2014. Fast Sparse Matrix-Vector Multiplication on GPUs for Graph Applications. In *Proc. HPC Netw., Stor. Anal. (SC '14)*. ACM, 781–792. DOI:http://dx.doi.org/10.1109/SC.2014.69

D. Barkai, K. J. M. Moriarty, and C. Rebbi. 1985. A Modified Conjugate Gradient Solver for Very Large Systems. *Comp. Phys. Comm.* 36, 1 (1985), 1 – 8. DOI:http://dx.doi.org/10.1016/0010-4655(85)90014-1

M. M. Baskaran and R. Bordawekar. 2008. Optimizing Sparse Matrix-Vector Multiplication on GPUs. *IBM RC24704* (2008).

N. Bell, S. Dalton, and L. Olson. 2012. Exposing Fine-Grained Parallelism in Algebraic Multigrid Methods. *SIAM J. Sci. Comp.* 34, 4 (2012), C123–C152. DOI:http://dx.doi.org/10.1137/110838844

N. Bell and M. Garland. 2009. Implementing Sparse Matrix-Vector Multiplication on Throughput-Oriented Processors. In *Proc. HPC Netw., Stor. Anal. (SC '09)*. ACM, Article 18, 11 pages. DOI:http://dx.doi.org/10.1145/1654059.1654078

A.T. Chronopoulos and C.W. Gear. 1989. s-step Iterative Methods for Symmetric Linear Systems. *Journal Comp. Appl. Math.* 25, 2 (1989), 153–168. DOI:http://dx.doi.org/10.1016/0377-0427(89)90045-9

M.M. Dehnavi, D.M. Fernandez, J. Gaudiot, and D.D. Giannacopoulos. 2013. Parallel Sparse Approximate Inverse Preconditioning on Graphic Processing Units. *IEEE Trans. Par. Dist. Sys.* 24, 9 (Sept 2013), 1852–1862. DOI:http://dx.doi.org/10.1109/TPDS.2012.286

J. Fang, A. L. Varbanescu, and H. Sips. 2011. A Comprehensive Performance Comparison of CUDA and OpenCL. In *Proc. Intl. Conf. Par. Proc.* 216–225. DOI:http://dx.doi.org/10.1109/ICPP.2011.45

R. Gandham, K. Esler, and Y. Zhang. 2014. A GPU Accelerated Aggregation Algebraic Multigrid Method. *Computers & Mathematics with Applications* 68, 10 (2014), 1151 – 1160. DOI:http://dx.doi.org/10.1016/j.camwa.2014.08.022

P. Ghysels and W. Vanroose. 2014. Hiding Global Synchronization Latency in the Preconditioned Conjugate Gradient Algorithm. *Par. Comp.* 40, 7 (2014), 224–238. DOI:http://dx.doi.org/10.1016/j.parco.2013.06.001

J. L. Greathouse and M. Daga. 2014. Efficient Sparse Matrix-Vector Multiplication on GPUs Using the CSR Storage Format. In *Proc. HPC Netw., Stor. Anal. (SC '14)*. ACM, 769–780. DOI:http://dx.doi.org/10.1109/SC.2014.68

M. J. Harvey and G. De Fabritiis. 2011. Swan: A Tool for Porting CUDA Programs to OpenCL. *Comp. Phys. Comm.* 182, 4 (2011), 1093–1099. DOI:http://dx.doi.org/10.1016/j.cpc.2010.12.052

M. R. Hestenes and E. Stiefel. 1952. Methods of Conjugate Gradients for Solving Linear Systems. *J. Res. Natl. Bureau of Standards* 49, 6 (1952), 409–436.

T. Jacques, L. Nicolas, and C. Vollaire. 1999. Electromagnetic Scattering with the Boundary Integral Method on MIMD Systems. In *High-Performance Computing and Networking*. LNCS, Vol. 1593. Springer, 1025–1031. DOI:http://dx.doi.org/10.1007/BFb0100663

K. Karimi, N. G. Dickson, and F. Hamze. 2010. A Performance Comparison of CUDA and OpenCL. *arXiv e-Print 1005.2581* (2010).

K. Kim and V. Eijkhout. 2013. Scheduling a Parallel Sparse Direct Solver to Multiple GPUs. In *IEEE Intl. Conf. Par. Dist. Sys. Workshops*. 1401–1408. DOI:http://dx.doi.org/10.1109/IPDPSW.2013.26

B. Krasnopolsky. 2010. The Reordered BiCGStab Method for Distributed Memory Computer Systems. *Procedia Comp. Sci.* 1, 1 (2010), 213–218. DOI:http://dx.doi.org/10.1016/j.procs.2010.04.024

M. Kreutzer, G. Hager, G. Wellein, H. Fehske, and A. R. Bishop. 2014. A Unified Sparse Matrix Data Format for Efficient General Sparse Matrix-Vector Multiply on Modern Processors with Wide SIMD Units. *SIAM J. Sci. Comp.* 36, 5 (2014), C401–C423. DOI:http://dx.doi.org/10.1137/130930352

V. W. Lee, C. Kim, J. Chhugani, M. Deisher, D. Kim, A. D. Nguyen, N. Satish, M. Smelyanskiy, S. Chennupaty, P. Hammarlund, R. Singhal, and P. Dubey. 2010. Debunking the 100X GPU vs. CPU Myth: An Evaluation of Throughput Computing on CPU and GPU. In *Proc. Intl Symp. Comp. Arch.* ACM, 451–460. DOI:http://dx.doi.org/10.1145/1816038.1816021

R. Li and Y. Saad. 2013. GPU-Accelerated Preconditioned Iterative Linear Solvers. *J. Supercomp.* 63, 2 (2013), 443–466. DOI:http://dx.doi.org/10.1007/s11227-012-0825-3

X. Liu, M. Smelyanskiy, E. Chow, and P. Dubey. 2013. Efficient Sparse Matrix-vector Multiplication on x86-based Many-core Processors. In *Proc. HPC Netw., Stor. Anal. (SC '13)*. ACM, 273–282. DOI:http://dx.doi.org/10.1145/2464996.2465013

M. Lukash, K. Rupp, and S. Selberherr. 2012. Sparse Approximate Inverse Preconditioners for Iterative Solvers on GPUs. In *Proc. HPC Symp.* SCS, Article 13, 8 pages.

G. Martinez, M. Gardner, and Wu chun Feng. 2011. CU2CL: A CUDA-to-OpenCL Translator for Multi- and Many-Core Architectures. In *IEEE Intl. Conf. Par. Dist. Sys.* 300–307. DOI:http://dx.doi.org/10.1109/ICPADS.2011.48

M. Naumov. 2012. Preconditioned Block-Iterative Methods on GPUs. *PAMM* 12, 1 (2012), 11–14. DOI:http://dx.doi.org/10.1002/pamm.201210004

J. Nickolls, I. Buck, M. Garland, and K. Skadron. 2008. Scalable Parallel Programming with CUDA. *Queue* 6, 2 (2008), 40–53. DOI:http://dx.doi.org/10.1145/1365490.1365500

C. Richter, S. Schops, and M. Clemens. 2014. GPU Acceleration of Algebraic Multigrid Preconditioners for Discrete Elliptic Field Problems. *IEEE Trans. Magn.* 50, 2 (Feb 2014), 461–464. DOI:http://dx.doi.org/10.1109/TMAG.2013.2283099

K. Rupp, Ph. Tillet, B. Smith, K.-T. Grasser, and A. Jüngel. 2013. A Note on the GPU Acceleration of Eigenvalue Computations. In *AIP Proc.*, Vol. 1558. 1536–1539.

Y. Saad. 1985. Practical Use of Polynomial Preconditionings for the Conjugate Gradient Method. *SIAM J. Sci. Stat. Comp.* 6, 4 (1985), 865–881. DOI:http://dx.doi.org/10.1137/0906059

Y. Saad. 2003. *Iterative Methods for Sparse Linear Systems, Second Edition*. SIAM. DOI:http://dx.doi.org/10.1137/1.9780898718003

Y. Saad and M. H. Schultz. 1986. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comp.* 7, 3 (1986), 856–869. DOI:http://dx.doi.org/10.1137/0907058

E. Saule, K. Kaya, and Ü. Catalyürek. 2014. Performance Evaluation of Sparse Matrix Multiplication Kernels on Intel Xeon Phi. In *Parallel Processing and Applied Mathematics*. Springer Berlin Heidelberg, 559–570. DOI:http://dx.doi.org/10.1007/978-3-642-55224-3_52

W. Sawyer, C. Vanini, G. Fourestey, and R. Popescu. 2012. SPAI Preconditioners for HPC Applications. *Proc. Appl. Math. Mech.* 12, 1 (2012), 651–652. DOI:http://dx.doi.org/10.1002/pamm.201210314

O. Schenk, M. Christen, and H. Burkhart. 2008. Algorithmic Performance Studies on Graphics Processing Units. *J. Par. Dist. Comp.* 68, 10 (2008), 1360–1369. DOI:http://dx.doi.org/10.1016/j.jpdc.2008.05.008

R. Strzodka and D. Göddeke. 2006. Pipelined Mixed Precision Algorithms on FPGAs for Fast and Accurate PDE Solvers from Low Precision Components. In *Proc. IEEE FCCM*. IEEE Computer Society, 259–270. DOI:http://dx.doi.org/10.1109/FCCM.2006.57

H. van der Vorst. 1992. Bi-CGSTAB: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comp.* 13, 2 (1992), 631–644. DOI:http://dx.doi.org/10.1137/0913035

M. Wagner, K. Rupp, and J. Weinbub. 2012. A Comparison of Algebraic Multigrid Preconditioners Using Graphics Processing Units and Multi-core Central Processing Units. In *Proc. HPC Symp.* SCS, Article 2, 8 pages.

H. F. Walker and L. Zhou. 1994. A Simpler GMRES. *Num. Lin. Alg. Appl.* 1, 6 (1994), 571–581. DOI:http://dx.doi.org/10.1002/nla.1680010605

I. Yamazaki, H. Anzt, S. Tomov, M. Hoemmen, and J. Dongarra. 2014. Improving the Performance of CA-GMRES on Multicores with Multiple GPUs. In *Proc. IEEE IPDPS*. IEEE Computer Society, 382–391. DOI:http://dx.doi.org/10.1109/IPDPS.2014.48

L.T. Yang and Richard P. Brent. 2002. The Improved BiCGStab Method for Large and Sparse Unsymmetric Linear Systems on Parallel Distributed Memory Architectures. In *Proc. Alg. Arch. Par. Proc.* 324–328. DOI:http://dx.doi.org/10.1109/ICAPP.2002.1173595

R. Yokota, J. P. Bardhan, M. G. Knepley, L. A. Barba, and T. Hamada. 2011. Biomolecular Electrostatics using a Fast Multipole BEM on up to 512 GPUs and a Billion Unknowns. *Comp. Phys. Comm.* 182, 6 (2011), 1272 – 1283. DOI:http://dx.doi.org/10.1016/j.cpc.2011.02.013